Complexity-Based Link Assignment for NASA's Deep Space Network for Follow-the-Sun Operations

Jigna Lad [1], Mark D. Johnston [2], Daniel Tran [3], David Brown [4], Kenneth Roffo [5] and Carlyn Lee [6]
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Road, Pasadena CA 91109 USA

NASA's Deep Space Network (DSN) recently underwent a paradigm shift in its operations approach called Follow the Sun Operations (FtSO) in an effort to increase efficiency for forthcoming expansion of the network. This change requires each Deep Space Communications Complex (DSCC) to remotely control the other two complexes' antennas during their local day shift, in contrast to locally controlling only their own antennas 24x7. Remote operations increases the workload of each complex during their day shift, specifically that of the Link Control Operators (LCOs), and presents a new challenge for planning and managing the distribution of responsibility for each link. A new DSN software assembly, the Link Complexity and Maintenance (LCM) software, was developed to support workload management for LCOs, as well as for planning site-local maintenance activities. The LCM deployment was a vital part of the transition to FtSO in November 2017. This paper discusses the architecture of LCM, its feature set, and lessons learned during its development and roll-out.

I. Introduction

NASA's Deep Space Network (DSN) [1], [2] provides communications and navigation support for all of NASA's planetary and deep space science mission, as well as several international partner missions. The DSN consists of three Deep Space Communications Complexes (DSCCs) evenly distributed in longitude at Goldstone, California; Madrid, Spain and Canberra, Australia. Currently, the network is comprised of 12 antennas; each complex contains one 70-meter antenna and three 34-meter antennas. There are approximately 40 frequent users of the DSN but this is expected to increase in the coming years with missions like InSight, the James Webb Space Telescope, Parker Solar Probe, and Mars 2020 launching in the next few years. Additionally, there is an expected increase in data rates, link complexity and the possibility of manned mission support. As a result, the DSN is expanding the number of available antennas to a total of 15 by the mid 2020's, and striving to automate operations as much as possible. To optimize the efficiency of operations in support of this increasing user community, the DSN underwent an operations paradigm shift in 2017. DSN changed its operations model such that each complex's operators not only control their own local antennas, but also remotely control those at the other two DSCC's during their local day shift. This paradigm change has been designated *Follow the Sun operations (FtSO)*[3].

FtSO eliminates the need for 24x7 staffing at each of the complexes, instead only requiring 9x7 staffing. The operational shift at each complex overlaps the next controlling station's shift by one hour to support a smooth handover of remote operations at the start and end of each shift. Although FtSO intends to improve the efficiency of DSN operations, it introduces many challenges to the Link Control Operators (LCOs) who operate and manage the antenna to spacecraft links. The workload of each complex increases from managing only four antennas to the full set of all 12 antennas. Additionally, the types of links have increased in complexity as the notion of Multiple Spacecraft per Antenna (MSPA) links has been upgraded from two simultaneous downlinks to four simultaneous downlinks. Other types of links also have increased complexity, due to spacecraft and ground system factors. These complexity factors will further increase in the coming years in addition to the growing user base, with cubesats

---

[1] Engineering Applications S/W Engineer, Data Services Group, Planning and Execution Systems Section
[2] Principal Scientist, Artificial Intelligence Group, Planning and Execution System Section
[3] Member of Technical Staff, Artificial Intelligence Group, Planning and Execution Systems Section
[4] Member of Technical Staff, Artificial Intelligence Group, Planning and Execution Systems Section
[5] Engineering Applications S/W Engineer, Modeling & Verification Group, Planning & Execution Systems Section
[6] Applications Software Engineer, Telecom Architectures Group, Communication Arch. and Research Section

expected to be a significant contributor. While automation of link operations has made great strides, there is still reliance on human-in-the-loop activities for each link, as well as interaction with the spacecraft Mission Operations Centers (MOC) when unplanned intervention is required during a link.

In order to address the challenges to LCOs brought on by Follow the Sun operations, new tools to plan and manage links were necessary. In mid 2017, the DSN team at NASA's Jet Propulsion Laboratory (JPL) began the deployment the Link Complexity and Maintenance tool (LCM) as a solution to some of these LCO challenges. LCM provides each complex with the ability to plan their shifts based on the current tracking schedule and on the complexity of the links contained in it, as well as to directly enter and manage other local DSCC activities such as maintenance and engineering. The primary focus of this paper is the development process, the available features of LCM, and the performance and impact of the tool since its deployment.

II. Software Overview

The Link Complexity and Maintenance software (LCM) was developed using agile development methodologies where continuous and frequent delivery and iterative requirement changes were welcomed. Requirements, tasks, and workflows were managed using ZenHub task boards which is a type of Kanban board that allows for tracking increments of work through stages as it is being done. Version control of the code base was achieved through the use of GitHub Enterprise, a commonly used source code management tool. In order to achieve an agile workflow, the end users of LCM, such as the link control operators (LCOs) and DSCC engineering staff, were kept in the loop throughout the development process to provide feedback and recommendations on features. Additionally, LCM was made available to DSCCs in multiple beta versions, to further augment the benefits of user in the loop development. This agile approach ensured that LCM was built to best supplement users' daily scheduling activities as well as the initial goal of alleviating the FtSO challenges.

The software itself was built on top of a prototype developed in an effort to characterize the complexity of link activities algorithmically and to evenly distribute the workload amongst LCOs based on constraints on their availability. This prototype laid the foundation for most of the front-end link assignment capabilities within the LCM web application. It also helped determine the factors that contribute to link complexity and how the complexity profile of a given link should be modeled[4]. One of the most useful aspects of the prototype was that it could be shown to LCOs for them to get a sense of how the software would work, and thus provide early feedback.

A. Architecture

LCM is built as a web application using NodeJS, a lightweight asynchronous, event-driven JavaScript runtime. Each complex locally runs their own instance of the LCM server to provide easy access to user networks within the DSCC's local area network (LAN). JPL hosts an additional read-only instance of the software that periodically synchronizes with each DSCC's data to provide an integrated view of operations across the entire network.

Each LCM server instance runs three NodeJS processes: the web application server, a database synchronizer, and an email notification engine. All three of these services are running continuously using a NodeJS package called *forever* for server instances hosted on MacOS and Linux, or using native Windows services for instances hosted on

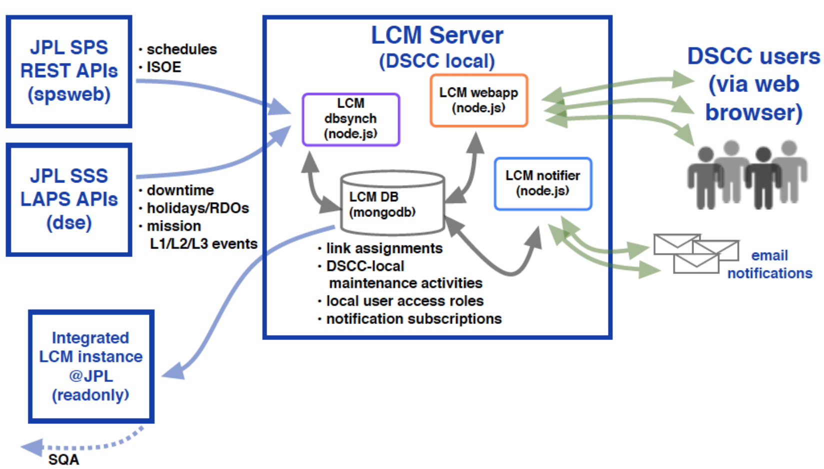

Figure 1. LCM Architecture

Windows operating systems. These programs act analogously to the POSIX command *nohup* by ensuring the processes do not stop, and if they are inadvertently terminated, they are respawned immediately. This is essential to ensuring LCM and its associated services are resilient. Lastly, each LCM server uses an instance of the NoSQL (non-relational) database, MongoDB, to store all pertinent data. The block diagram in Figure 1 further illustrates the architecture of LCM showing both the internal configuration as well as its external interfaces.

B. Database Synchronization and Link Creation

In order to keep LCM data up to date with the most recent schedule revisions, LCM incorporates a database synchronization process that runs every 10 minutes. This synchronizer queries a representational state transfer-based (RESTful) application programming interface (API) that provides detailed information of the latest version of the DSN tracking schedule. Tracks contain all information relevant to generating a link which includes the spacecraft or other user, time bounds, asset used, and references to other grouped activities such as MSPAs or arrays. During database synchronization, links are generated from a set of tracks and are assigned an overall *complexity class* based on the contained tracks; these link objects will later be assigned to an operator through the LCM user interface.

The complexity class of each link is determined by the highest class of all tracks contained within it. Complexity is assigned as follows:

Link Type	Complexity Class Value
Downlink Only	1
No Ranging	2
Ranging	3
Delta Differential One-Way Ranging (DDOR)	4

The higher the complexity value, the more the demand on the LCO's attention. The link creation process also generates a complexity profile that provides a discretized complexity value for every 5 minute interval over the duration of the link. In general, links will have a "spiked" complexity at the start of track setup (Start of Activity, or SOA), and again at the start of actual tracking (Beginning of Track, or BOT), as these periods requires more work and attention from the operator. For the remainder of the track duration the complexity profile is modeled as exponentially decreasing until the start of the teardown interval between End of Track (EOT) and End of Activity (EOA). The teardown interval has a lower complexity value, but is also modeled as an exponential[4]. The overall complexity of the link is the maximum of all the time-phased complexity values of the contained tracks, thus allowing for the variations due to spacecraft swaps for MSPAs where up to four tracks are running at once as part of a single link.

In addition to complexity values based on track attributes, which are specified in the schedule far in advance, LCM also has the capability to look at detailed sequences of events (SOEs) which are sometimes not fully specified until days ahead of track execution. The synchronication process fetches SOEs for tracks occurring within 48 hours in the future, and scans them for factors that may impact complexity. As a specific example, Mars missions that are occulted by the planet will undergo periodic loss and re-acquisistion of signal. These times generally require special attention by the LCO to verify that all is performing as expected, and so LCM includes a complexity increment in the overall profile for tracks with these events.

This characterization of complexity is essential to the link assignment process as it allows LCO supervisors to fairly distribute the workload amongst operators on shift, accounting for the fact that complexity changes from the start of a track as time passes after the link is built and the signal is acquired. Additionally, this allows LCO's to only be assigned to a level of activity for which they are prepared. For example, a new operator may not yet be trained to handle certain types of links, and so their maximum complexity can be limited to avoid assigning them links with too high a rating.

LCM also provides users with information about scheduled asset downtimes and critical mission events to give context to the schedule. This data is synchronized from a RESTful API web service provided by a separate tool, the Loading Analysis and Planning Software (LAPS) managed out of JPL and used for long range planning[5]. Downtimes are also used in LCM to define key intervals for scheduling specific maintenance and engineering activities – another key feature of LCM.

Lastly, the database synchronizer performs an import/export procedure in order to push data local to each DSCC to the integrated JPL instance of LCM. Every hour the DSCC LCM servers perform an export of key database content, essentially snapshotting the state at that time. This snapshot is then compressed and uploaded to a file system portal hosted at JPL. The JPL instance queries this portal hourly to pull the exported data and populate the JPL instance with the aggregated data from all three complexes. As a result, the JPL instance is able to show link assignments and maintenance activities for all three complexes in a read-only view, at most an hour out of date.

C. Notifications

LCM data is volatile; it is frequently changing from schedule updates through database synchronization, and due to user activity within the system. It is often pertinent that users are informed immediately of changes that directly affect activities scheduled to them. As a result, an email notification system was integrated with LCM. Through a preference menu in the user interface, users can subscribe to several types of notifications. These include changes to link assignments, introduction or changes to link or maintenance conflicts, as well as alerts for upcoming activities. The email notification process works by performing a scan of the database audit table every five minutes. The scans check for changes to system data based on user/system activity since the last scan. These changes are aggregated and sent to users based on their subscriptions. Notifications also provide a mechanism for LCO admin users to be alerted that there are unassigned links approaching, prompting them to take immediate action. Lastly, LCM system administrators can also subscribe to a system-type notification indicating that there has been failure to synchronize the database – either because of a general system failure or an inability to reach the RESTful API web services at JPL.

D. Web Application

The web application is implemented as a classical client-server setup with routing middleware. The web server handles all access to the MongoDB data and the client provides an interactive website for users to work with that data. The NodeJS middleware, Express, is used to route requests from the client to the server and handle authorization. The database is protected, such that users must possess the right permissions in order to view and edit any data and the middleware performs this authorization check before routing requests to the server. The web application also has a built-in auditing system. All user activity and system activity (such as database synchronization) are logged to the database audit table. This table allows for traceability when system problems arise, as well as useful data for capturing LCM metrics. The following sections go into detail on the various views and features exposed to users through the web application client.

III. User Interface

The LCM client is served as a single page application, where isolated features are grouped into five views that can be toggled via tabs Figure 2). Above the tabs, in the interface header, are several interactive buttons, menus and indicators for modifications to the view. The header includes access to view modification buttons which allow a user

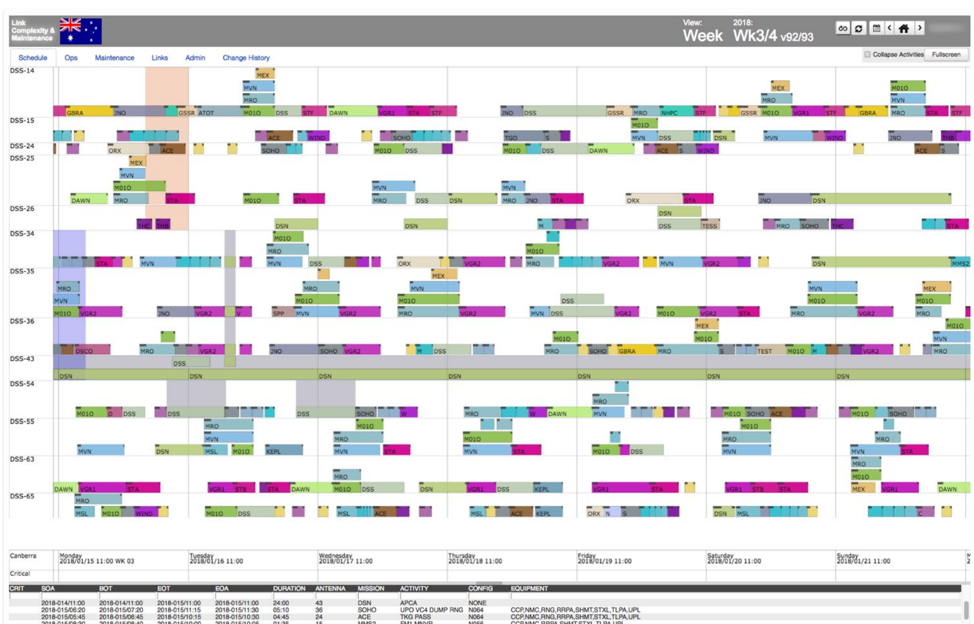

Figure 2. LCM Web Application – Schedule View for a single week

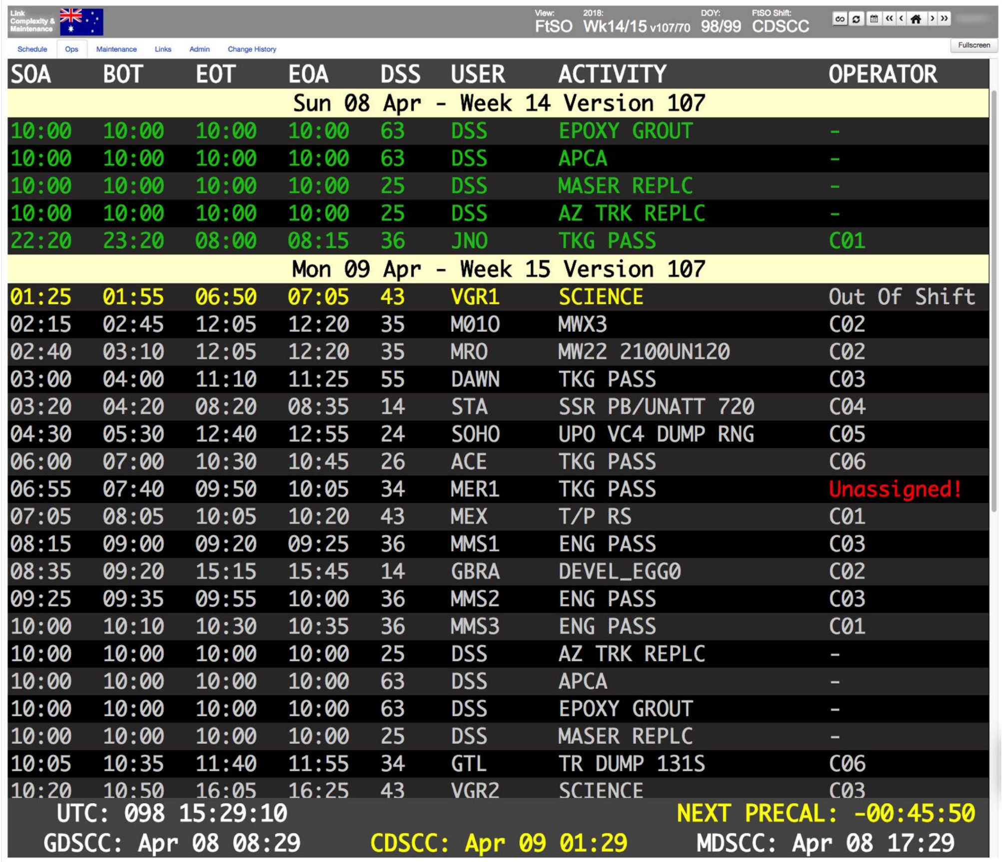

SOA	BOT	EOT	EOA	DSS	USER	ACTIVITY	OPERATOR
Sun 08 Apr - Week 14 Version 107							
10:00	10:00	10:00	10:00	63	DSS	EPOXY GROUT	-
10:00	10:00	10:00	10:00	63	DSS	APCA	-
10:00	10:00	10:00	10:00	25	DSS	MASER REPLC	-
10:00	10:00	10:00	10:00	25	DSS	AZ TRK REPLC	-
22:20	23:20	08:00	08:15	36	JNO	TKG PASS	C01
Mon 09 Apr - Week 15 Version 107							
01:25	01:55	06:50	07:05	43	VGR1	SCIENCE	Out Of Shift
02:15	02:45	12:05	12:20	35	M010	MWX3	C02
02:40	03:10	12:05	12:20	35	MRO	MW22 2100UN120	C02
03:00	04:00	11:10	11:25	55	DAWN	TKG PASS	C03
03:20	04:20	08:20	08:35	14	STA	SSR PB/UNATT 720	C04
04:30	05:30	12:40	12:55	24	SOHO	UPO VC4 DUMP RNG	C05
06:00	07:00	10:30	10:45	26	ACE	TKG PASS	C06
06:55	07:40	09:50	10:05	34	MER1	TKG PASS	Unassigned!
07:05	08:05	10:05	10:20	43	MEX	T/P RS	C01
08:15	09:00	09:20	09:25	36	MMS1	ENG PASS	C03
08:35	09:20	15:15	15:45	14	GBRA	DEVEL_EGG0	C02
09:25	09:35	09:55	10:00	36	MMS2	ENG PASS	C03
10:00	10:10	10:30	10:35	36	MMS3	ENG PASS	C01
10:00	10:00	10:00	10:00	25	DSS	AZ TRK REPLC	-
10:00	10:00	10:00	10:00	63	DSS	APCA	-
10:00	10:00	10:00	10:00	63	DSS	EPOXY GROUT	-
10:00	10:00	10:00	10:00	25	DSS	MASER REPLC	-
10:05	10:35	11:40	11:55	34	GTL	TR DUMP 131S	C06
10:20	10:50	16:05	16:25	43	VGR2	SCIENCE	C03

Figure 3. LCM Ops View: a tabular view of the near-term schedule, highlighting link assignment status

to jump to a specific time range or jump to the next or last view. There is also a user menu where the user can set preferences, such as their default view type, email subscriptions and more. A refresh icon that allows a user to load the latest update of the client data from the server automatically turns red when changes to the current view are detected. This immediately lets the user know that they are viewing stale data, but without disrupting what they are doing. There is also an automatic refresh setting that will perform the update whenever stale data is detected. Within each of the tabs, users are exposed to different features and the following sections go into detail on each of those.

A. Summary Views

LCM provides users with two summary views of schedule data, the Schedule view and the Ops view. Figure 2. shows a screen shot of the user interface default schedule view. This tab is viewable by all users of LCM and gives an at-a-glance view of DSN activities within a given time frame in a timeline. Each of the colored bars represent a DSN track labeled by its user. Areas where bars overlap each other on the same antenna band often compose a multiple spacecraft per antenna (MSPA) link. The schedule view also draws various colors of background shading to indicate antenna downtime (gray), complex holidays (orange), and critical mission events (red, green, yellow). Users can customize this view to zoom in or out of the time range in view, toggle background shading, collapse overlapping activities, filter antennas in view, and spotlight on specific missions or equipment usage.

The second view available to all users in LCM is the Ops tab shown in Figure 3. This view displays all links chronologically in tabular form and is often on display on large monitors at the DSCCs as a general-purpose activity monitor. The view gives a good overview of links currently underway and those upcoming as it updates with countdown timers in real time. It also gives a good indication of approaching links that are unassigned or might need immediate attention with the last column indicating operator assignment. Unassigned links are highlighted as shown in red in the "Operator" column, and links out of shift are identified as such.

B. Link Assignment

Link Assignment is performed within the Links tab of LCM (Figure 4). Here, there are two classes of users – admins and general users. Both of these role types can access the tab and be assigned to shifts and link activities. The link admins have the ability to assign users with link roles to the shift in view. This essentially means that the operator can be assigned links during the shift and is expected to support any assignments made to them. Link assignment is achieved by moving the link bars off of the antenna timelines and onto the operator timelines. Assigned links turn gray on the antenna timeline, so it is easy to spot any that remain unassigned. Assignments can be done manually or through an auto-assign feature.

Auto-assign will algorithmically assign links to operators while attempting to minimize conflicts and to even out the workload. First, auto-assign will try to evenly distribute links amongst operators on shift, respecting the maximum complexity and maximum number of links that are allowed for each operator. Then it will make a number of iterations through the assignments, trying to move links in conflict to another operator until as many as possible are removed. After the maximum number of iterations is reached, there are at most a small number of changes that LCOs might want to perform to finalize the assignments. These manual changes could be based on personal preference or other factors not implemented into the algorithm.

Additionally, there is a feature to interactively assign all links in a shift that are on a given antenna to an individual operator. This is a desirable assignment configuration as the operator has to bring up and configure displays for each link, and there are fewer changes for links on the same antenna.

C. Link Conflicts and Constraints

Each LCO has two attributes that determine the type and amount of links that they can handle – maximum link complexity and maximum number of links. These attributes impose a constraint on link assignment. If a link is assigned that exceeds an operator's maximum link complexity value at any point during the link duration as defined by the its complexity profile, LCM will indicate a conflict for the operator. Similarly, if an operator is assigned more overlapping links than their maximum allowed number, LCM will indicate a violation. Violations are shown in the link assignment tab with red background shading during the intervals in conflict.

On top of these operator-based constraints, LCM also has some other built in conditions for indicating conflicts on link assignment. First, an operator should not be assigned additional links during an 3- or 4-spacecraft segment of an MSPA link. This rule has been implemented because MSPAs are usually more demanding for an operator with multiple setup and teardown intervals to support, and are generally managed separately from any other link. The

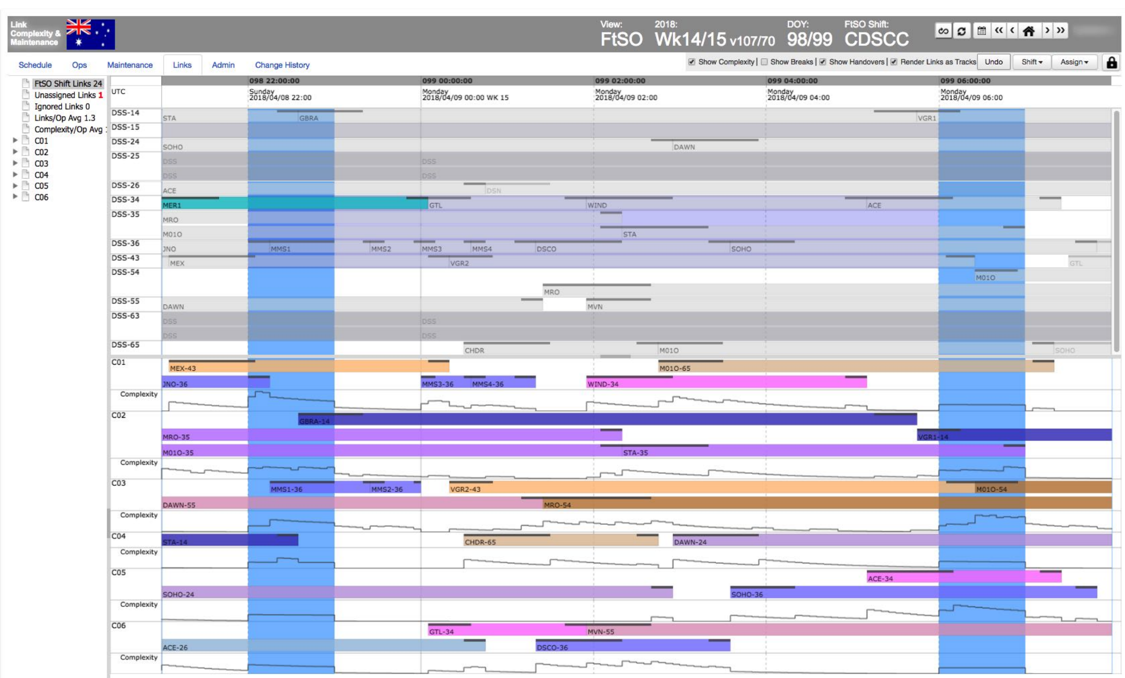

Figure 4. LCM Link Assignment for a single shift: the schedule is at the top, and operator timelines are below. Assigned links are grayed out. Here, one link remains unassigned (MER1 on DSS-34).

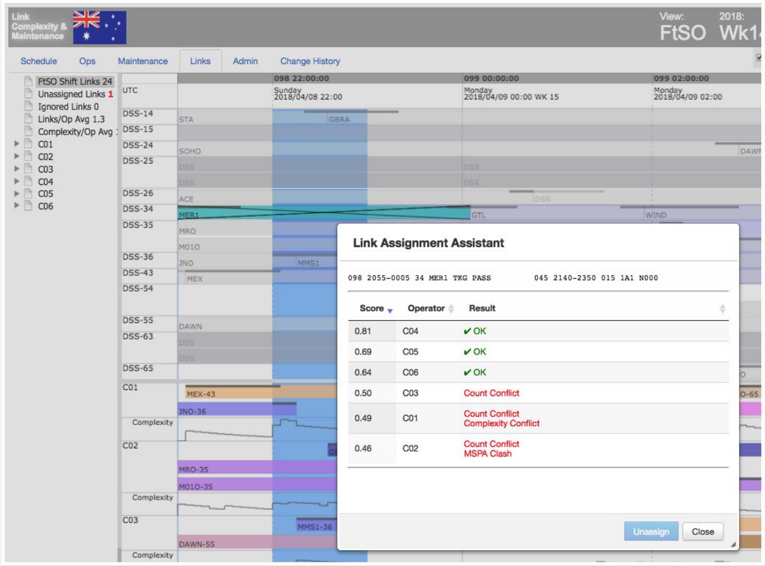

Figure 5: The link assignment assistant illustrating options for assignment of an unassigned track. Valid options are scored, and infeasible choices have a clear explanation of the conflict.

next rule is that an operator should not be assigned multiple links with the same start of activity (SOA) time. SOA is typically the highest point of the complexity profile, and the time when briefings from the spacecraft MOC occur, so an operator's attention should not be divided at that time. The last special conflict condition is specific to the Magnetospheric Multiscale (MMS) project, a four-spacecraft formation flying mission. Consecutive MMS tracks are represented in LCM as one link even though there are four physical spacecraft, and, as a result, each MMS link has multiple setup and teardown intervals to transition from one spacecraft to the next. These intervals are notably shorter than standard transition times (15 minutes vs the normal 45 minutes). LCM indicates a violation if there are any overlaps with other assigned links near these MMS short transition intervals. Specifically, if there are overlapping SOAs during the transition or 15 minutes before the transition, there will be a violation and similarly, any BOT or EOT overlapping a transition is a violation

During the link assignment process, it can be difficult to resolve conflicts by just manually moving links around from operator to operator. To help with this process, LCM has an assignment assistant feature (Figure 5) that guides users through conflict resolution. When activated for a given link, the assignment assist feature will show the result of hypothetically moving that link to each one of the operator timelines – whether moving the link will cause the other operator a violation, what type of violation(s), or whether the move will not cause any issues. There is also a score calculated related to workload, so that a higher score indicates a move that will least impact that operator. Figure 5 shows a sample dialog where moving the selected link has three valid options as well as three others that would cause conflicts. This feature is useful when auto-assign is unable to fully resolve link conflicts, or when and additional manual changes are necessary.

Users may want to try some "what-if "scenarios in link assignment for a subset of links while still being able to make use of the auto-assign feature. To enable this, any link can be locked in place. When a link is locked, it cannot be moved manually or through automatic assignment and as a result, its assignment stays fixed while all other unlocked links can be moved around. Another feature that provides a similar capability is the Do-Not-Assign option on a link. This applies to links not yet assigned to an operator and allows users to exclude it from automatic assignment. Do-Not-Assign is useful when a link is known not to need LCO attention (e.g. it will complete during the handover interval, or is a special engineering activity), and should not keep showing up as "unassigned".

The last feature of the link assignment tab is the ability to schedule breaks for an operator. LCOs generally try to stagger their lunch times and breaks to ensure that there is always an operator available to support links, even during normal lunch hours. To keep track of this, each operator timeline has an associated break timeline where customizable break intervals can be defined. LCOs performing link assignment can use these as avoidance intervals for the operator to ensure that they can get a break during their shift time. This is also useful for tracking and avoidance of an operator's irregular absences such as meetings, training sessions, etc.

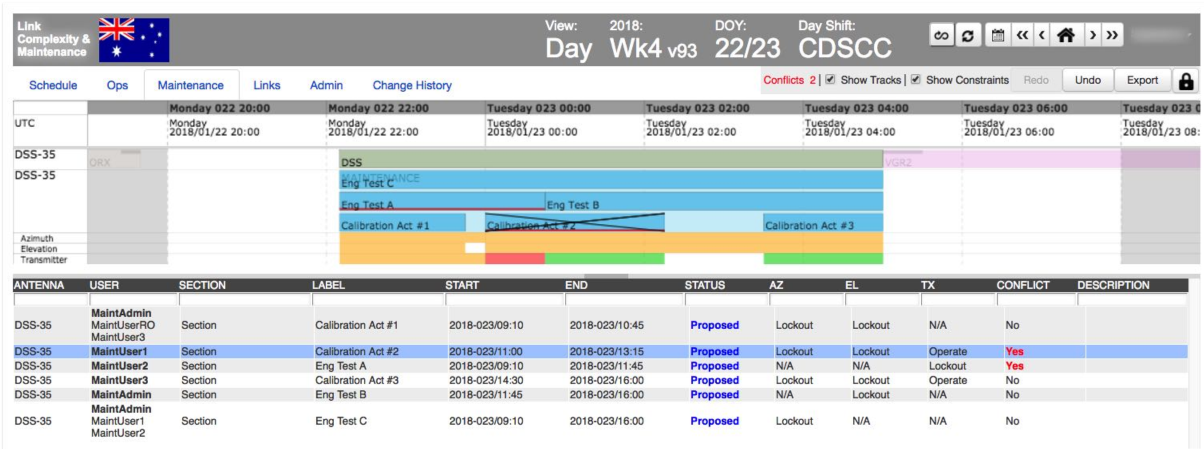

Figure 6: LCM maintenance scheduling

D. Maintenance Scheduling

The second primary feature of LCM is for maintenance activity scheduling. DSN assets frequently undergo scheduled maintenance such as testing, upgrades and calibrations as well as routine health and safety inspections. These activities must occur during the downtime intervals defined in the LAPS tool and synchronized over to the LCM database. Although these downtime intervals are predefined, the maintenance staff do have flexibility in scheduling more granular work items within them and can control the assignment of personnel to those tasks; the maintenance tab of LCM allows users to do so.

Within the maintenance tab, there are three classes of users - maintenance admins, maintenance users and read-only maintenance users. Each of the maintenance admins and users have the ability to create and edit activities within the predefined maintenance intervals. Maintenance activities have several customizable attributes- the start time, end time, point of contact (POC), users, resource reservations, and status. The difference between maintenance admin and regular maintenance users are that admins can edit all activities, whereas regular users can only edit those that they own. Additionally, admins have elevated privileges on modifying the status of an activity. Maintenance read only members can be assigned to any maintenance activity but have no write permission to maintenance data.

E. Maintenance Workflow

Maintenance activities go through a workflow represented by the status field. Newly created activities are always in a PROPOSED state. From here, a maintenance admin can change the state to APPROVED or DENIED. From APPROVED, the status can be changed to CANCELED. DENIED and CANCELED activities can be edited by their creator, or any other admin users, and changes will automatically return the status back to PROPOSED. This workflow is imposed to enforce the validity of maintenance activities in the schedule. Each status type has an associated color, clearly distinguished by the color of the maintenance activity bar in the timeline view.

F. Maintenance Constraints

Similar to constraints on link assignments, LCM also provides conflict detection on maintenance activities based on maintenance constraints. Maintenance conflict can occur for a number of reasons and are indicated as red background shading behind intervals in conflict. Most trivially, a conflict will be indicated on an activity that falls outside of a predefined downtime interval; in other words, maintenance must be fully contained within a scheduled downtime. The second type of maintenance activity conflict results from a resource conflict. As mentioned earlier, resource reservations can be made on a maintenance activity. This reservation can be operate, lockout or none for the antenna's azimuth, elevation or transmitter resources. Because it is valid to have as many activities as desired on a single antenna, conflicts are flagged if any of those reserve the same resource with both operate and lockout overlapping any time interval. Each antenna timeline has three sub-timelines, one for each of the reservations to provide a visual representation of the resource usage and of any conflicting time intervals. In the example in Figure 6, two example activities are conflicting in their requested use of the transmitter resource.

IV. LCM Administration

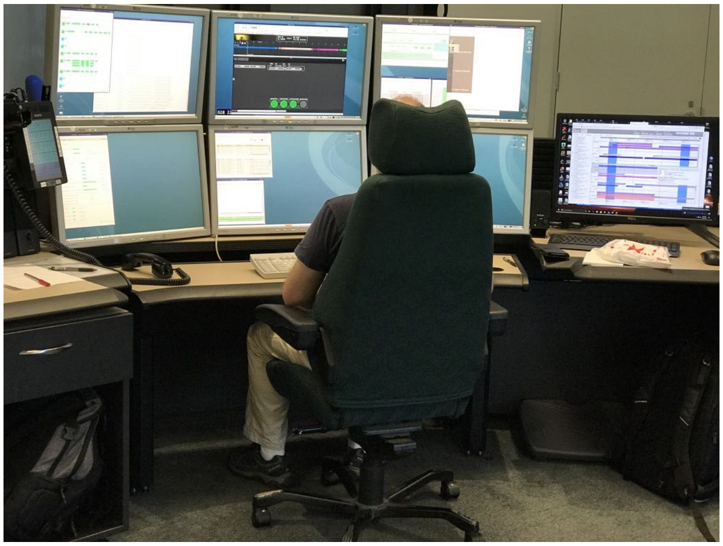

Figure 7: LCM in use at the Canberra Deep Space Communications Complex on the screen at the right.

The LCM software is hosted locally at each DSCC, therefore most of the system administration is handled directly by local IT staff. JPL assisted the IT staff through initial deployment of the tool and provided the necessary training and tutorial material for LCM's kickoff.

User administration within the web application is handled by LCM sysadmin users – those with sysadmin role. As previously mentioned, there are different classes of users available in LCM: sysadmin, ops, opsadmin, maintadmin, maint, maintro. The last five were explained in the link assignment and maintenance activity sections. All of the features available to the other five user classes are also available to sysadmin users, and additionally, they have several abilities within the user interface that the other roles do not. A sysadmin user has access to an Admin tab where modifications to LCM users and their attributes can be made as well as changes to some system properties applicable to the LCM server configuration. An admin can create new users, make modifications to their contact information, team assignments, and permissions (which user classes they are part of), as well as the operator attributes specifying maximum complexity level and maximum number of links. This information is all stored within the LCM database and also synchronized to the JPL instance. This way, users can log in, with their same credentials, to the integrated instance.

LCM provides a comprehensive auditing mechanism for tracking activity within the tool. As stated in the section about notifications, all system and user activity are recorded to a database audit table. This information is available to users in a change log form through the user interface. The change log displays changes for specific links and maintenance activities to show the attributes that have changed with their before and after values. The change log will also indicate the time and author of the change. This is useful for administrative purposes when problems arise and troubleshooting is required.

In order to comply with JPL's Service Quality Assessment (SQA) standards, LCM has also provided a SQA API into the integrated instance. This API returns a snapshot of LCM data, specifically of link assignments and user data, for any given complex at the start or end of their shift. The API allows for querying data for any SQA snapshot, and is used to derive metrics that can be correlated with other activities of the DSN.

V. Conclusion

LCM has been in use for over five months to date and has been enthusiastically received as a tremendous asset for staff at each of the DSCCs. It has enabled the following capabilities:

- the ability to view up to date DSN schedule data with accurate track data, sequences of events, critical mission events and downtimes, all in a single contextual view
- the ability to plan and manage the workload of Link Control Operator who are supporting links on all 12 DSN antennas in Follow the Sun (Remote) Operations
- the ability to schedule and distribute the workload of local maintenance and engineering activities
- the ability to receive automated email notifications when schedule changes or link/maintenance assignment changes occur

- management of LCM user information such as team and shift assignments and constraints on link complexity and maximum number of links
- localized system administration of the LCM software and server

Users of LCM have been able to utilize these new capabilities in their daily work. Because LCM runs on a separate LAN from the operations monitor and control workstations, it can be kept up side by side without consuming screen space on the former (Figure 7). The software has been performing as expected with only a few minor issues discovered and fixed in point releases. After initial deployment, a user survey was conducted to identify and prioritize outstanding feature requests, which will be the subject of future updates.

Acknowledgments

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[1] "Deep Space Network - NASA Jet Propulsion Laboratory." [Online]. Available: https://deepspace.jpl.nasa.gov/. [Accessed: 04-Apr-2018].

[2] W. A. Imbriale, *Large Antennas of the Deep Space Network*. Wiley, 2003.

[3] M. D. Johnston, M. Levesque, S. Malhotra, D. Tran, R. Verma, and S. Zendejas, "NASA Deep Space Network: Automation Improvements in the Follow-the-Sun Era," presented at the IJCAI Workshop on AI in Space, Buenos Aires, Argentina, 2016.

[4] D. Tran and M. D. Johnston, "Automated Operator Link Assignment Scheduling for NASA's Deep Space Network," presented at the 9th International Workshop on Planning and Scheduling for Space, Buenos Aires, Argentina, 2015.

[5] M. D. Johnston and J. Lad, "Integrated Planning and Scheduling for NASA's Deep Space Network – from Forecasting to Real-time," presented at the SpaceOps, Marseilles, France, 2018.